EXHIBIT 18

FILED UNDER SEAL

JEDEC

Global Standards for the Microelectronics Industry

DDR3 High Density DIMMs and 3DS

Tom Kinsley
Micron Technology

Server Memory Forum 2011

Micron

Agenda

- What is 3DS?
- 3DS Functional Highlights
- Potential applications

Note: This document provides a high level overview and does not include all aspects that are reflected in the device datasheet. Before operation of a 3DS device, please refer to a complete datasheet.

What is 3DS Technology

3DS Technology

- Years ago we stacked TSOP "packages"

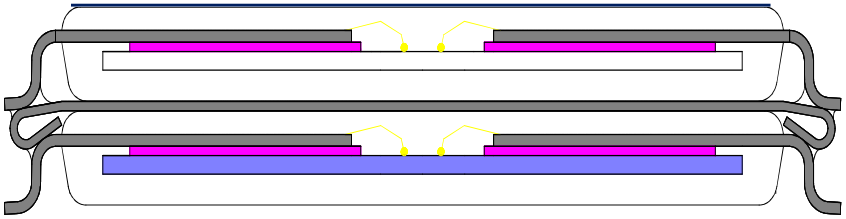

- As the speeds got faster and the die got smaller we migrated to die stacking.

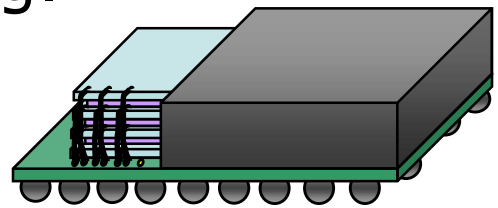

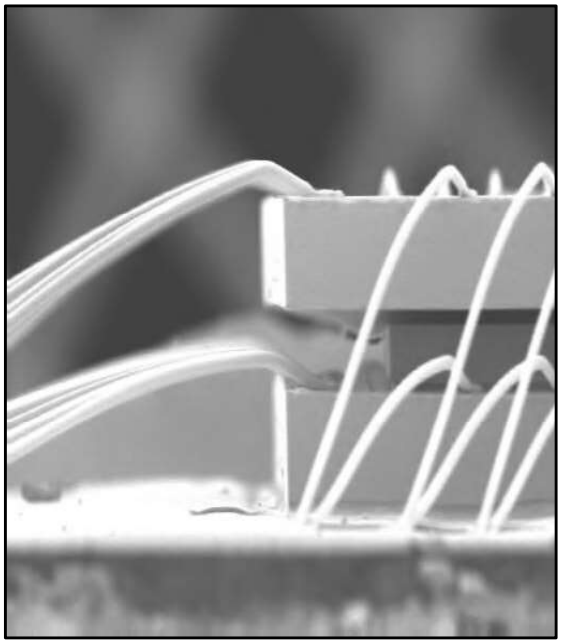

3DS Technology

- But there are limitations of the traditional die stack
  - Signal Integrity
    - High capacitive loads (QDP pin can approach ~12pf)
    - High inductance die interconnect

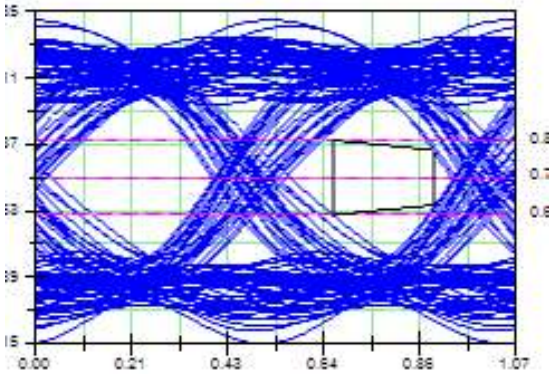

  - Mechanical
    - Thick packages may restrict system socket pitch and package thermals

- <u>Bottom line</u> – traditional stacking is becoming obsolete

3DS Technology

- **3DS** (Three Dimensional Stack) is more than a new method to stack die, 3DS is an innovative improvement in DRAM technology...
  - Most 3DS devices will utilize TSVs interconnects

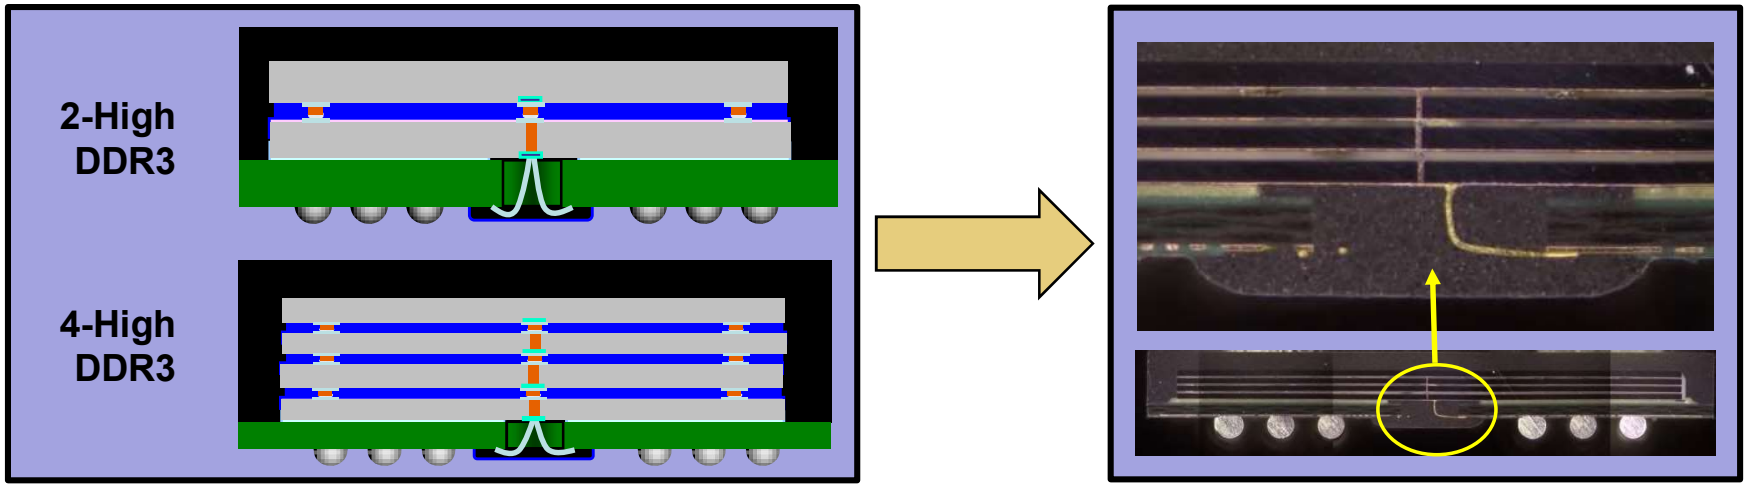

3DS Technology

- Through Wafer Interconnects (TSVs) allow a greater number of die to die interconnections which can operate at faster speeds and with a lower pj/bit. This allows the DRAM designer to incorporate a Master/Slave type architecture with many improvements over traditional DDR3 stacks.

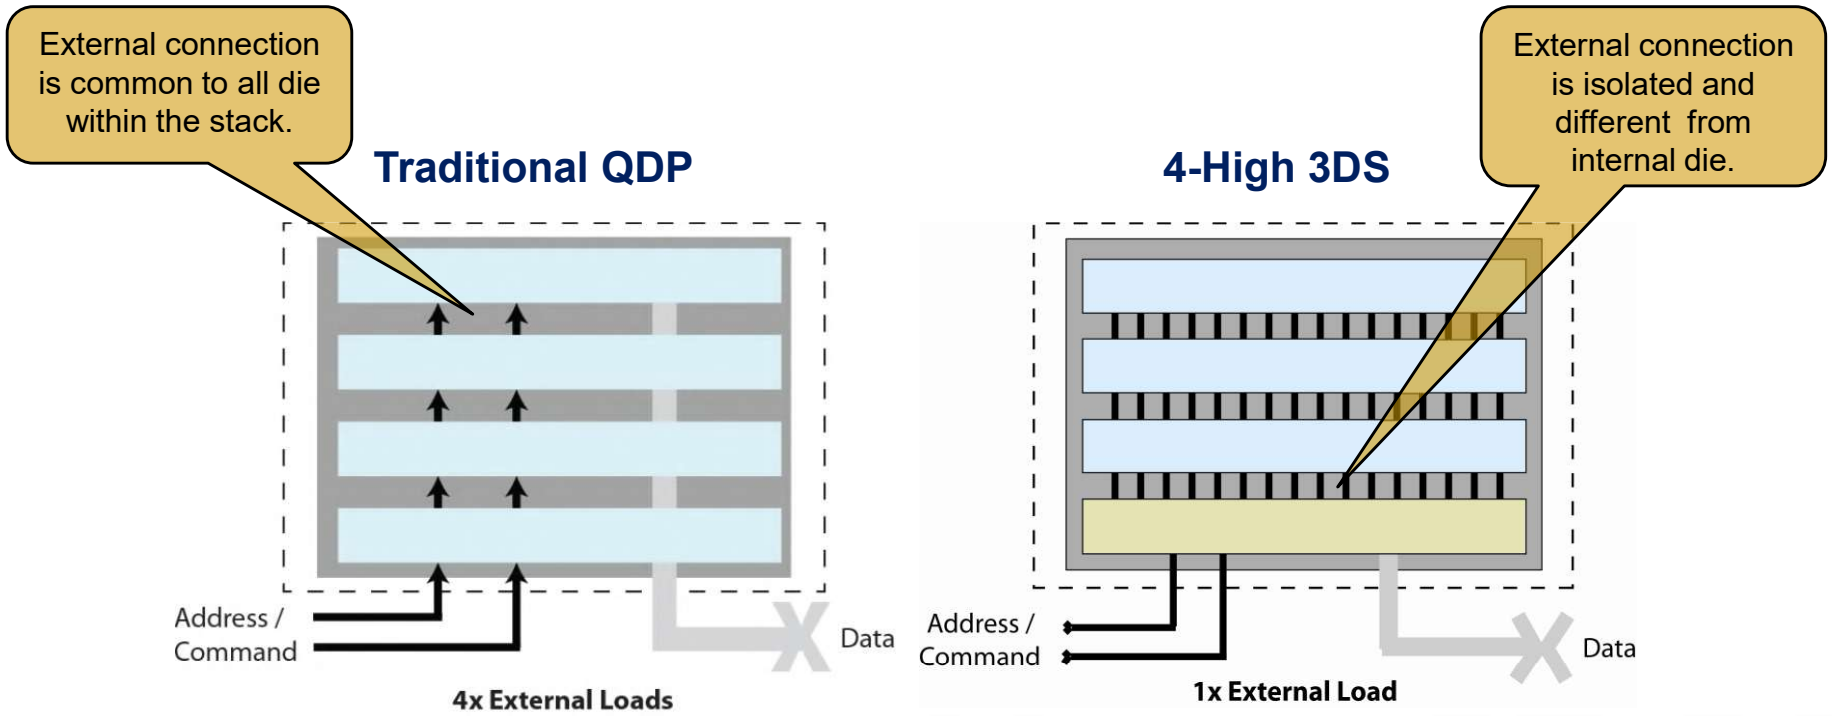

3DS Technology

- Key Design Improvements include:
  - Single electrical load for each ball
    - All inputs and external I/O are common to all ranks and are buffered
  - Only one DLL per stack
    - All ranks share same DLL
  - Reduced active logic
    - Some redundancy may be removed
  - Improved interconnect interface
    - Die to die electrical interface may not reflect standard I/O
  - Low voltage DDR3L support
    - Optimized for Vdd=VddQ=1.35V levels only

3DS Technology

- DDR3 3DS Benefits:
  - Lower Power
    - Device power
    - External power (allows weaker drivers and optimized termination due to reduced loading)
  - Better Thermals
    - Due to lower power and thinner packages
  - Higher Efficiency
    - Improved Rank to Rank access timing
  - Better Performance
    - Single loads promote faster bus speeds, better signal integrity
  - Less overhead
    - Single set of mode registers
    - Only one ZQ pin and requires less calibration time
    - Single CKE promotes easier power down enter/exits

3DS Functional Highlights

3DS Functional Highlights

- **DDR3 vs. DDR3 3DS**
  - The 3DS DDR3 specification utilizes the base (JESD79-3 DDR3 SDRAM) specification, but there are a few differences between the base DDR3 functionality and DDR3 3DS functionality which should be identified.
  - The following pages highlight the minor differences between the base and 3DS SDRAM.

3DS Functional Highlights

- **Ball-out**
  - The 3DS device uses a similar ball-out as the standard DDR3 DDP/QDP stacks with the following exceptions

VSS	VDD	NC
VSS	VSSQ	DQ0
VDDQ	DQ2	DQS
VSSQ	DQ6	DQS_n
VREFDQ	VDDQ	DQ4
NC	VSS	RAS_n
ODT	VDD	CAS_n
CS1_n	CS0_n	WE_n
VSS	BA0	BA2
CS2_n[3], NC	A3	A0
CS3_n[4], NC	A5	A2
VDD	A7	A9
VSS	RESET_n	A13

NU/TDQS_n	VSS	VDD
DM/TDQS	VSSQ	VDDQ
DQ1	DQ3	VSSQ
VDD	VSS	VSSQ
DQ7	DQ5	VDDQ
CK_t	VSS	CID[1], NC
CK_c	VDD	CKE
A10/AP	ZQ	NC
A15[2], NC	VREFCA	VSS
A12/BC_n	BA1	VDD
A1	A4	VSS
A11	A6	VDD
A14	A8	VSS

Only a single ODT input

CID is new

Only one (1) CKE pad

Only one (1) ZQ pad

1.Only used for 3DS devices with eight logical ranks
2.Only used for 4Gb and 8Gb logical rank densities
3.Only used for 3DS devices with four or eight logical ranks
4.Only used for 3DS devices with four or eight logical ranks

3DS Functional Highlights

- **Logical block diagram**
  - 3DS may support 2-High, 4-High or 8-High as x4 or x8
  - Note there is only one each (CKE and ODT inputs) per stack

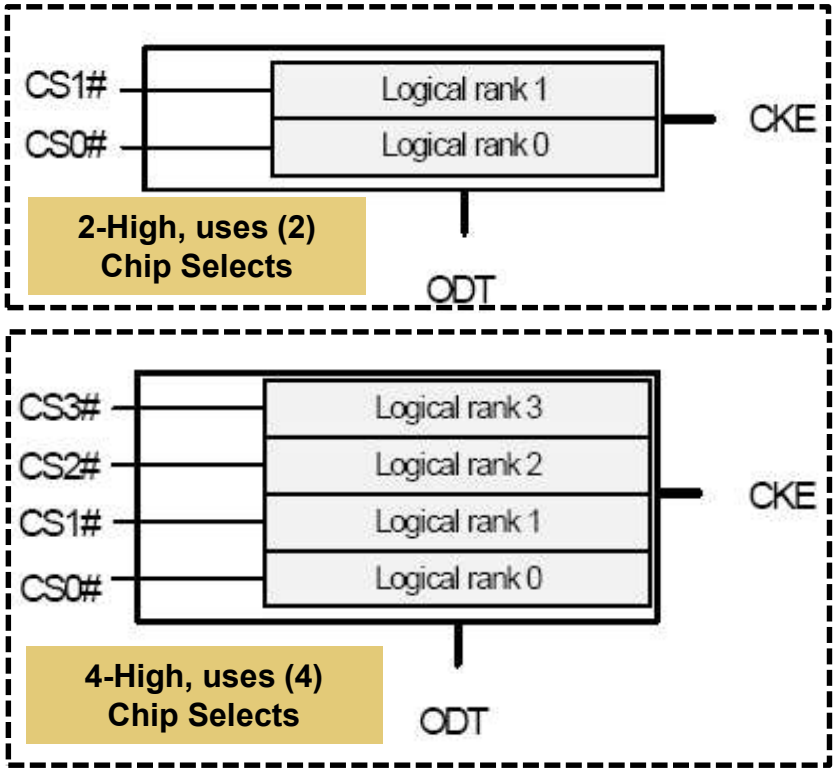

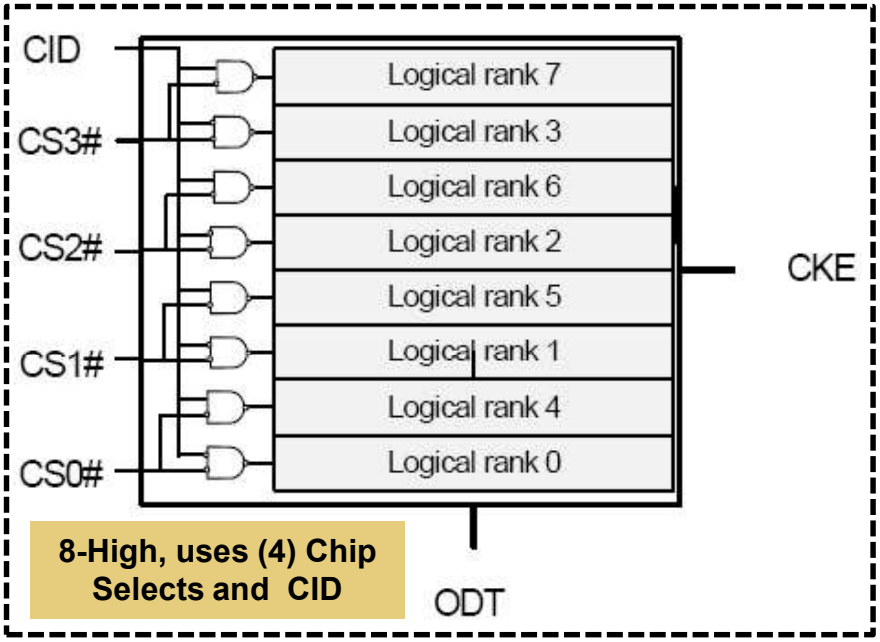

3DS Functional Highlights

- **Mode Registers**
  - Like planar DDR3 SDRAMs, DDR3 3D Stacked SDRAMs have four Mode Registers.
  - One set of registers controls the entire stack regardless if the 3DS stack has two, four or eight logical ranks, and they must be programmed via a Mode Register Set (MRS) command.
    - Note PRE commands (or PRE commands to each open bank) have to be issued to all logical ranks that may have open banks before the 3DS device can successfully execute a MRS command.
    - MRS command will only be executed if CS0# is low.

3DS Functional Highlights

- **Active Commands**
  - In a 3D Stacked DDR3 SDRAM the four chip select pins and the CID pin select the logical rank. No more than one logical rank ACTIVE command can be initiated simultaneously to DDR3 3DS devices, i.e. no more than a single chip select can be active when an ACTIVE command is send to a 3DS device.
  - 3DS devices introduce a few new timing parameters
    - Minimum time between ACTIVE commands
      - To the Same Logical Rank is **tRRD_slr**
      - To Different Logical Ranks is **tRRD_dlr**
    - Minimum time for the the first Active command after four have already been issued
      - To the Same Logical Rank is **tFAW_slr**
      - To Different Logical Ranks is **tFAW_dlr**

3DS Functional Highlights

- **Active Commands**
  - Sample timing diagrams showing t_{RRD} and t_{FAW} for both same logical rank and different logical ranks.

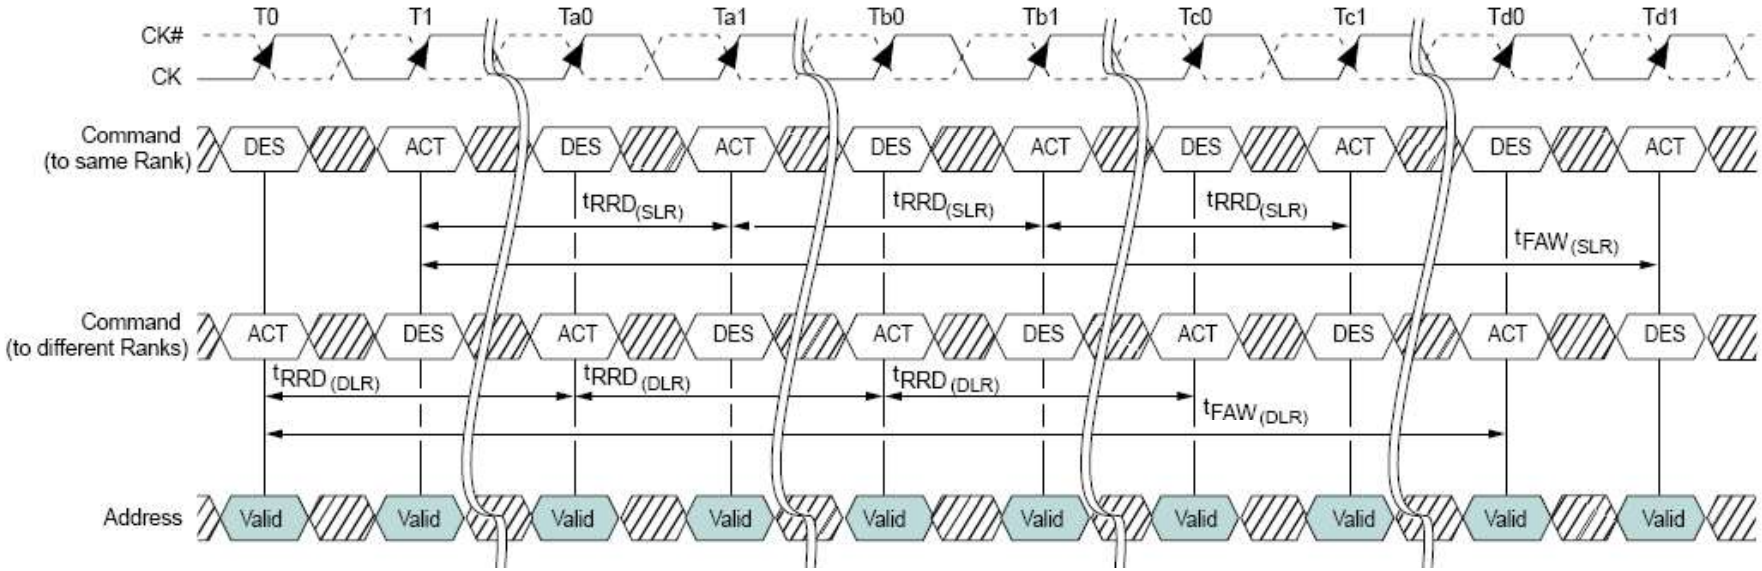

3DS Functional Highlights

- **Precharge and Precharge All Commands**
  - The Single Bank Precharge (PRE) and Precharge All Banks (PREA) commands apply only to selected logical ranks of a 3D Stacked SDRAM. It is illegal to issue PREA commands to multiple logical ranks simultaneously if they have more than 8 open rows in total.
  - PRE commands (or PRE commands to each open bank) have to be issued to all logical ranks with open banks before the device can enter Self Refresh mode.
  - 3D Stacked SDRAMs have the same values for tRP, tRTP, tRAS and tDAL as planar DDR3 SDRAMs of the same frequency.

3DS Functional Highlights

- **Read and Write Commands**
  - For 3DS devices the minimum time from issuing successive read-to-read or write-to-write commands to different logical ranks is t_{CCD} even to different logical ranks. No additional latency needs to be added for logical rank to logical rank turnaround time.
    - For back-to-back read-to-write or write-to-read commands issued to different logical ranks, turnaround times are completely identical as to same logical rank.
  - READ Latency for the 3DS device is equal to the CAS Latency of the base device plus one or more clock cycles, depending on speed bin.
    - Example if CL=7 clocks for a DDR3-1066 monolithic device, it will be CL+1 (8 clocks) or CL+2 (9 clocks) for the 3DS DDR3-1066 device.

3DS Functional Highlights

- **Read Commands**
  - Timing example showing contiguous READ data when accessing two different logical ranks.

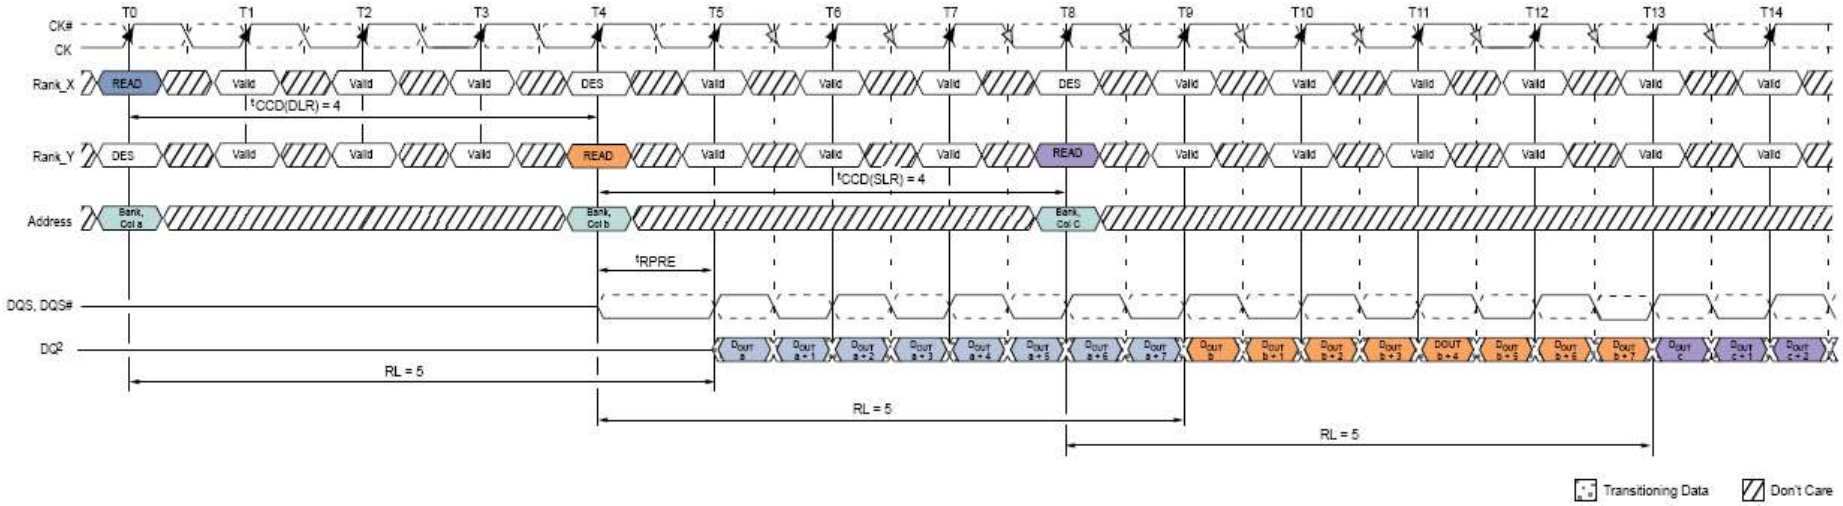

3DS Functional Highlights

- **Refresh Command**
  - No more than one logical rank Refresh Command can be initiated simultaneously to DDR3 3D Stacked SDRAMs, i.e. no more than a single chip select can be active when a Refresh command is sent to a 3D stacked device.
  - The minimum refresh cycle time to a single logical rank (=tRFC_slr) has the same value as tRFC for a planar DDR3 SDRAM of the same density as the logical rank.
  - The minimum time between issuing refresh commands to different logical ranks is specified as tRFC_dlr. After a Refresh command to a logical rank, other valid commands can be issued before tRFC_dlr to the other logical ranks that are not the target of the refresh.

3DS Functional Highlights

- **Refresh Command**
  - Sample timing diagrams showing same and different logical rank examples

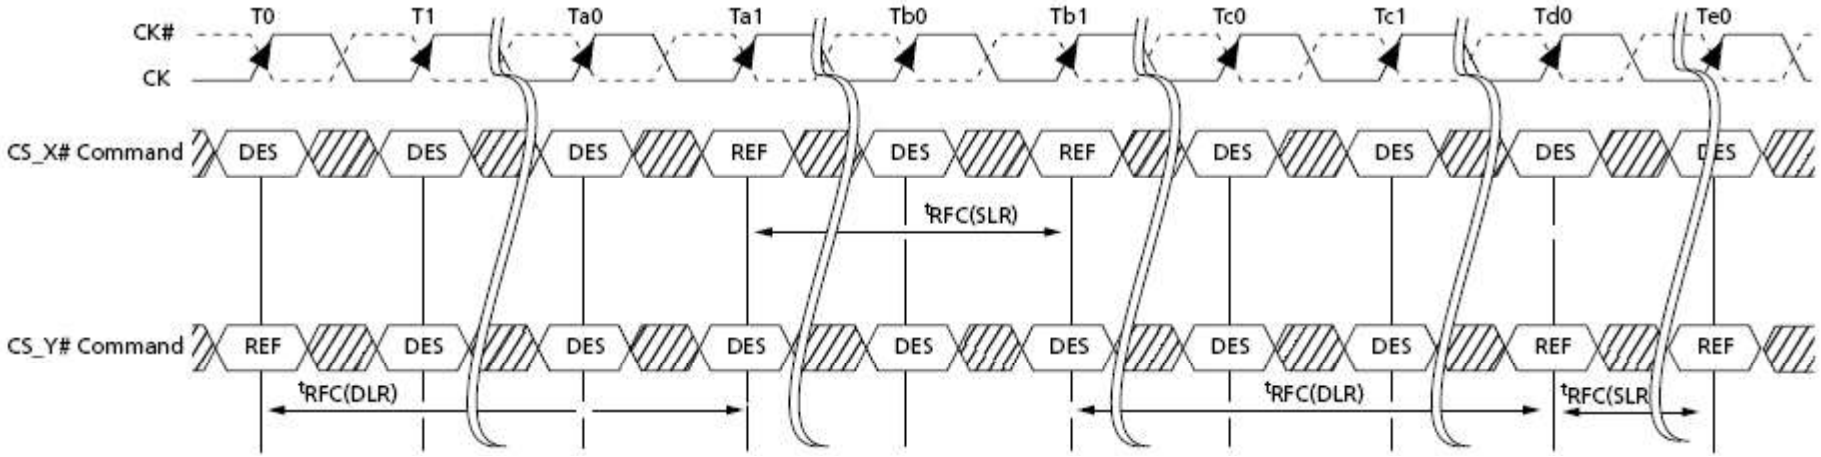

3DS Functional Highlights

- **Self-Refresh Operation and Power-Down Modes**
  - Since there is only one CKE pin per 3DS device, all logical ranks enter self refresh and power down together.
    - CS0# is required to be low to enter SR mode.
  - Self-Refresh exit (SRX) and power-down exit (PDX) apply to all logical ranks in a 3D Stacked device and is caused by the Low-to-High transition of the single CKE pin.

3DS Functional Highlights

- **ZQ Calibration Commands**
  - Each 3DS package will have a single ZQ calibration pin, independent of the number of logical ranks in the stack.
  - The lowest numbered chip select (CS0_n) of each package should be associated with the master logical rank. The ZQ pin is associated with master logical rank.
  - The calibration procedure and the result should adhere to JEDEC DDR3 component specification. The host may issue ZQ calibration command to each logical rank. The SDRAM can choose to ignore the ZQ commands to the non-master logical rank or repeat the calibration of the I/O attached to the master die.

Potential server applications

Potential server applications

- In the RDIMM channel loading of the DQ signals is typically the limiting factor for channel speed.

Example waveforms are at DDR3-1066 with Vdd=VddQ=1.5V

  - A QR module has four DQ loads per slot when using conventional DDR3 stacks. For this reason it is very difficult to get good signal integrity in a two slot system at the lower DDR3L voltage and while operating at higher clock frequencies.

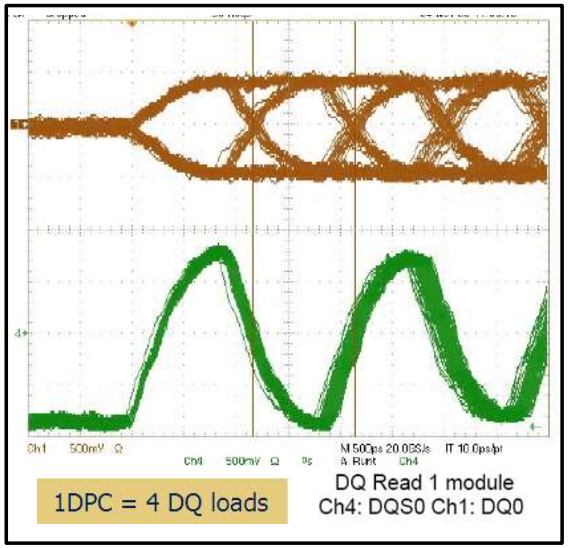

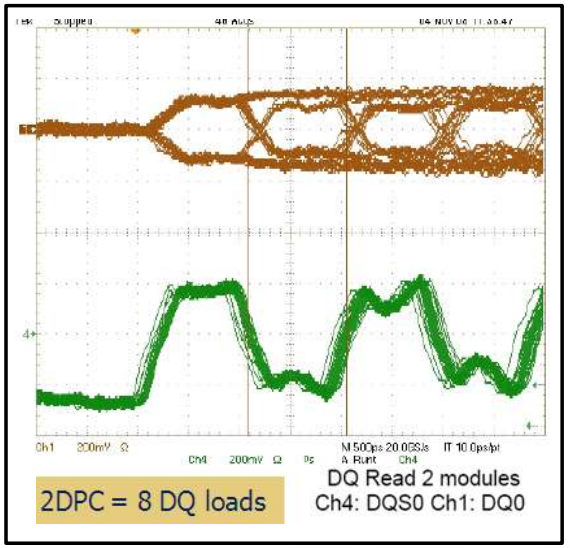

Potential server applications

- As with READs, the WRITE signals start to collapse as well. In addition to the high *capacitive load, there is typically termination turned on at one of more DRAM in the channel.

Example waveforms are at DDR3-1066 with Vdd=VddQ=1.5Vc

* At 1DPC there could be up to12pF on each DQ node, with 2DPC that could increase to 24pF (C_{IO}=1.5pf to 3.0pf depending speed grade).

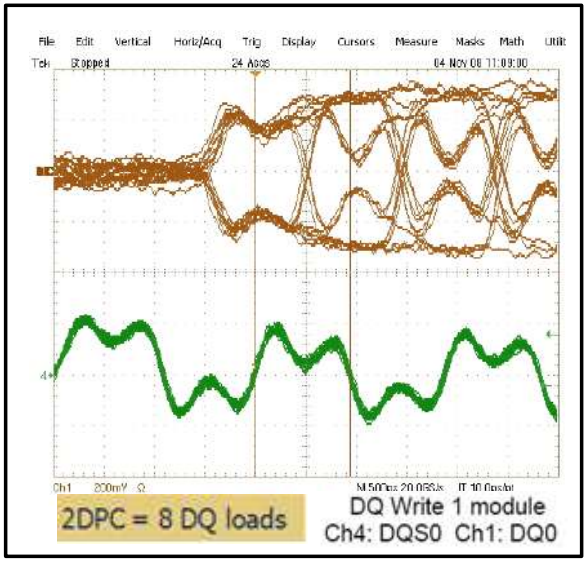

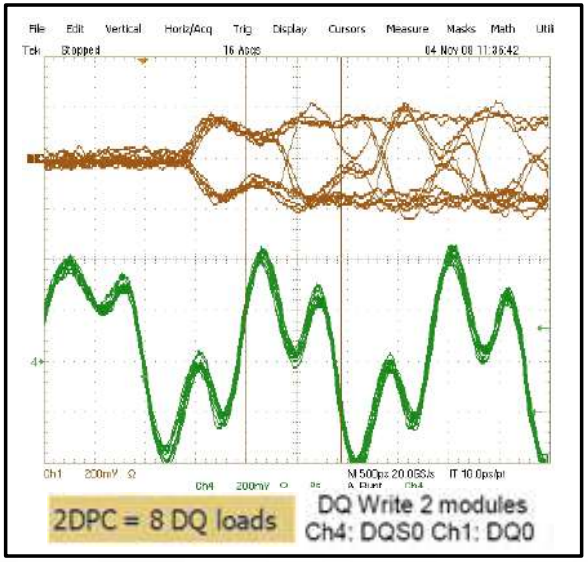

Potential server applications

- Building the RDIMM using 3DS yields two obvious solutions.
  - Option 1)
    - One row of 4-high 3DS devices which reflect 1 DQ load per slot (3pf max).

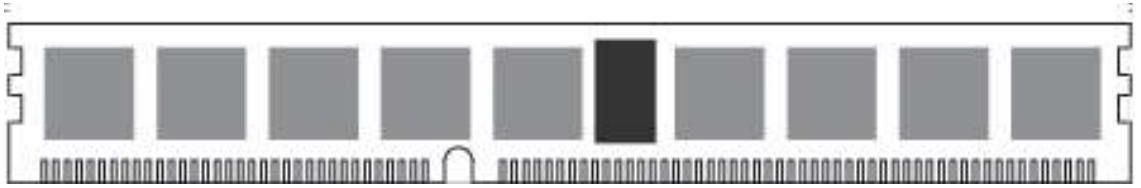

  - Option 2)
    - Two rows of 2-High 3DS devices which reflect 2 DQ loads per slot (6pf max).

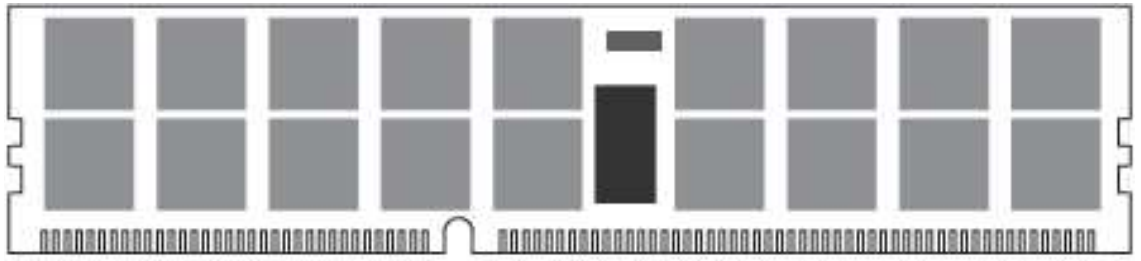

Potential server applications

- Using a like channel (speed, voltage and slots populated) with a single load 3DS RDIMM, the data eye opens up enough that there is actually timing margin on the DQ signals.

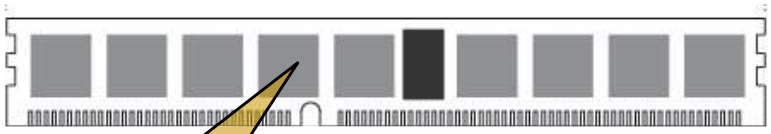

Each stack is 4-High 3DS with a single DQ load

READs

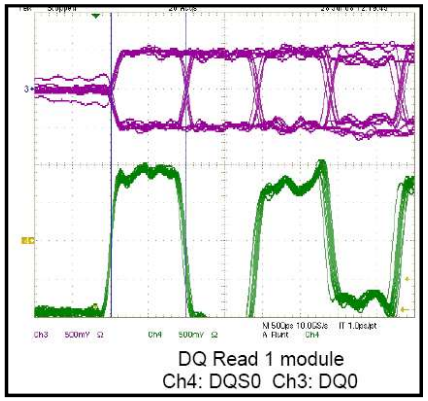

DQ Read 1 module
Ch4: DQS0 Ch3: DQ0

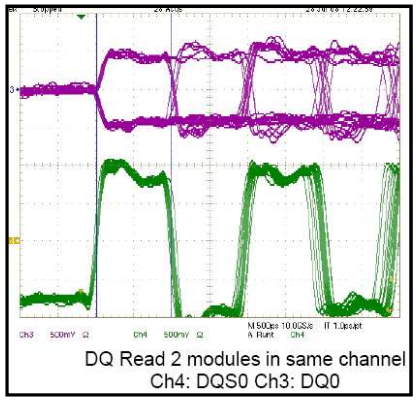

DQ Read 2 modules in same channel
Ch4: DQS0 Ch3: DQ0

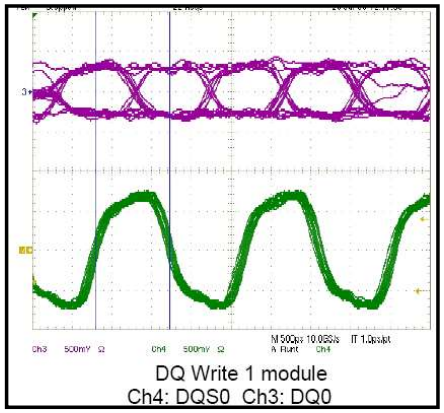

DQ Write 1 module
Ch4: DQS0 Ch3: DQ0

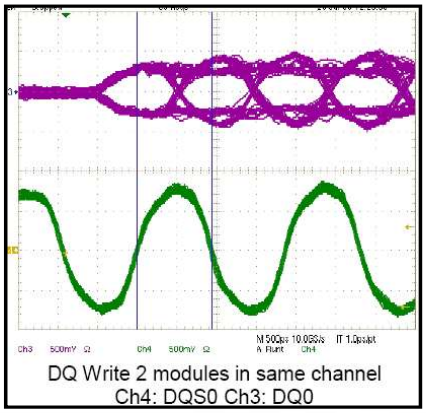

DQ Write 2 modules in same channel
Ch4: DQS0 Ch3: DQ0

WRITEs

Potential server applications

- Using a like channel (speed, voltage and slots populated) with a dual load 3DS RDIMM, the data eye is better than standard RDIMM.

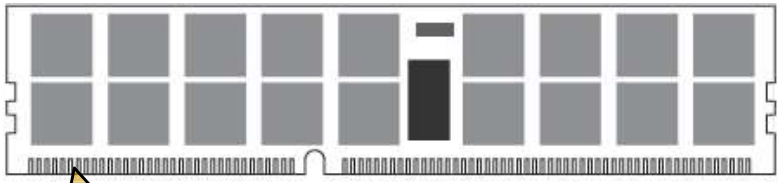

Each stack has a single DQ load, but there are two stacks on each node so the channel see two DQ loads.

READs

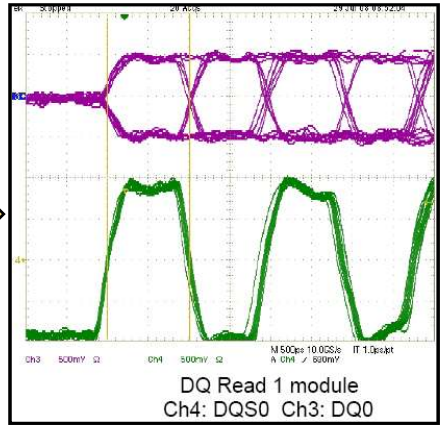

DQ Read 1 module
Ch4: DQS0 Ch3: DQ0

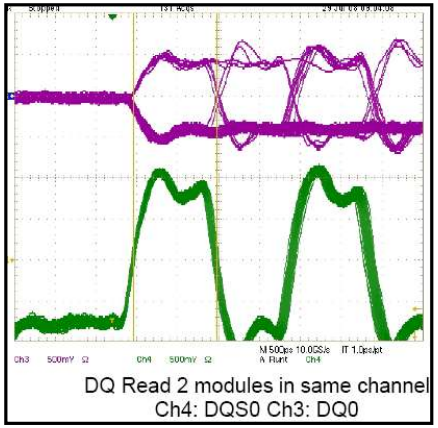

DQ Read 2 modules in same channel
Ch4: DQS0 Ch3: DQ0

WRITEs

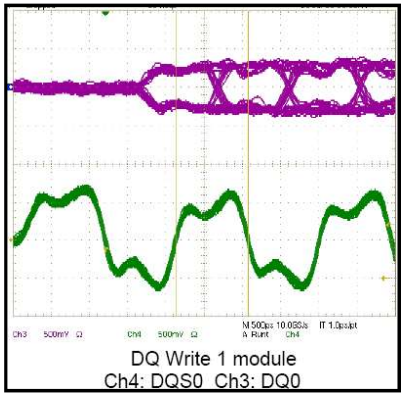

DQ Write 1 module
Ch4: DQS0 Ch3: DQ0

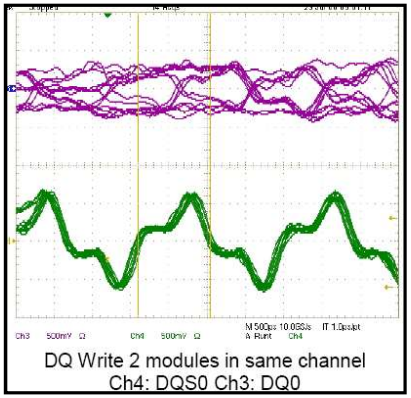

DQ Write 2 modules in same channel
Ch4: DQS0 Ch3: DQ0

Potential server applications

- 3DS RDIMM are possible, but they are not expected to become widely accepted because:
  - Single Row (single load) QR RDIMM requires 4-high stacks.
    - Any module built with 4-High memory stacks typically reflects a much higher cost than a like density module built with planar or 2-high stacked memory.
  - Most server channels support three (3) sockets, but most channels only support eight (8) chip selects per channel. As such a QR RDIMM is limited to only two slots per channel. There may be other ways to get the same density for less cost if utilizing all three sockets.

Potential server applications

- LRDIMMs (Load Reduced DIMMs) overcome the problem of limited chip selects in the channel as LRDIMM support "Rank Multiplication". With rank multiplication a LRDIMM can support up to eight ranks per slot for multiple slots, *even if there are only eight chip selects per channel.

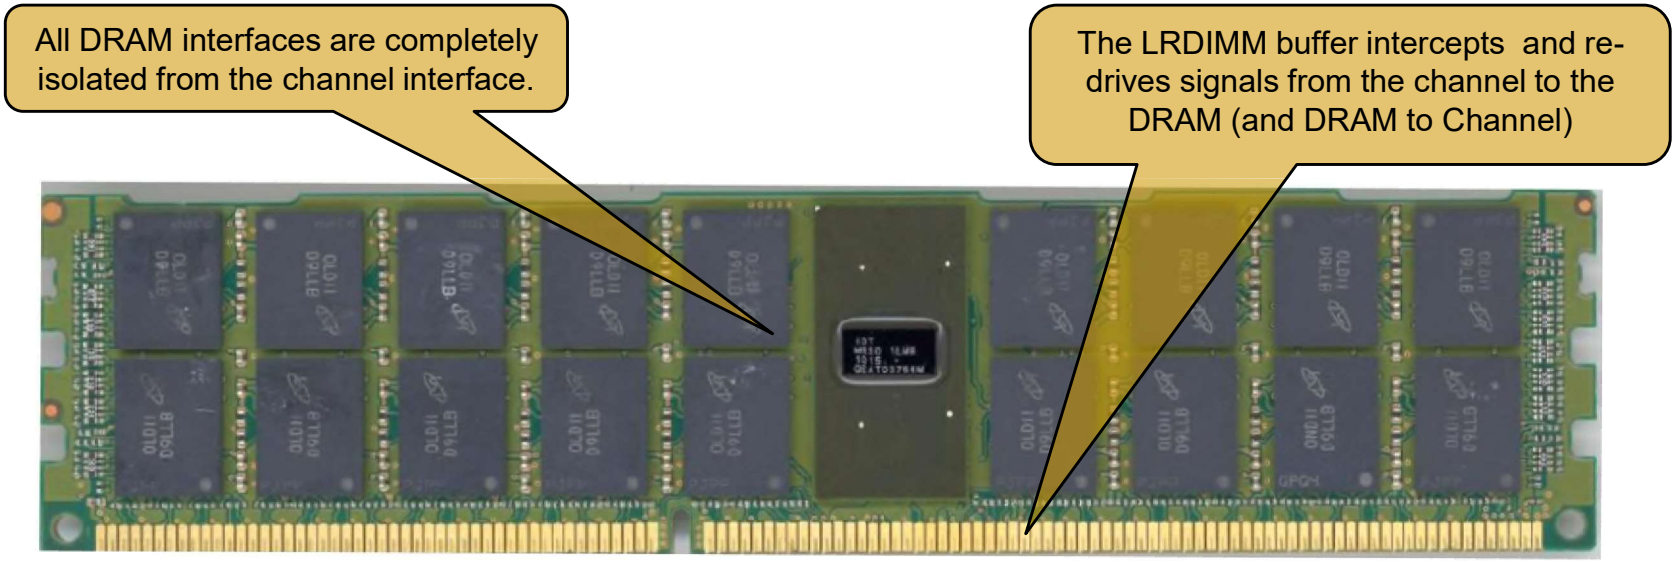

**LRDIMM incorporate higher order channel addresses to perform rank multiplication, the amount of ranks per slot that can be supported depends on the number of unused addresses available.*

Potential server applications

- LRDIMMs build with DDP or QDP can have high loading on the address and command lines which may hamper the data eye on the post buffer side.
  - For Example: (8R x4) LRDIMM built with QDP conventional stacks.

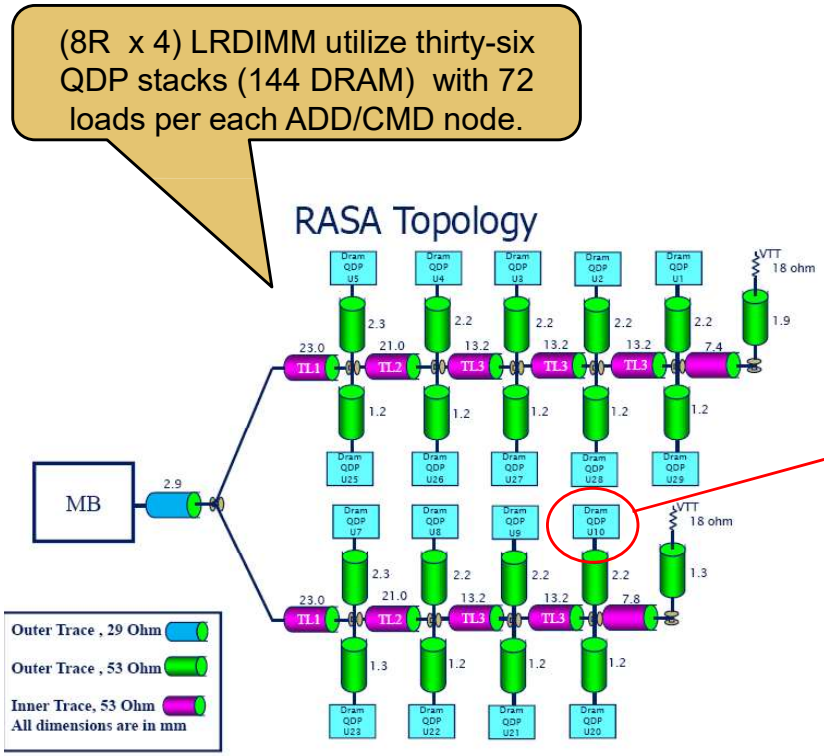

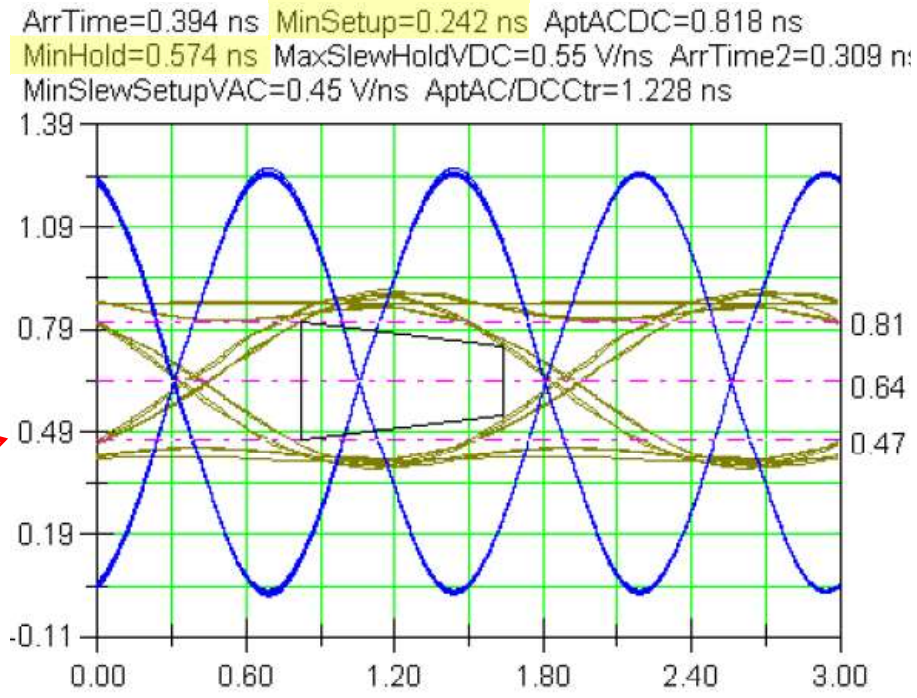

(8R x4) LRDIMM, U10, DDR3-1333, @ 1.35V, slow corner

Potential server applications

- Building the module with 3DS devices, the loading is reduced from four loads per package stack to one load per package stack which allows a good data eye even at DDR3L-1600.

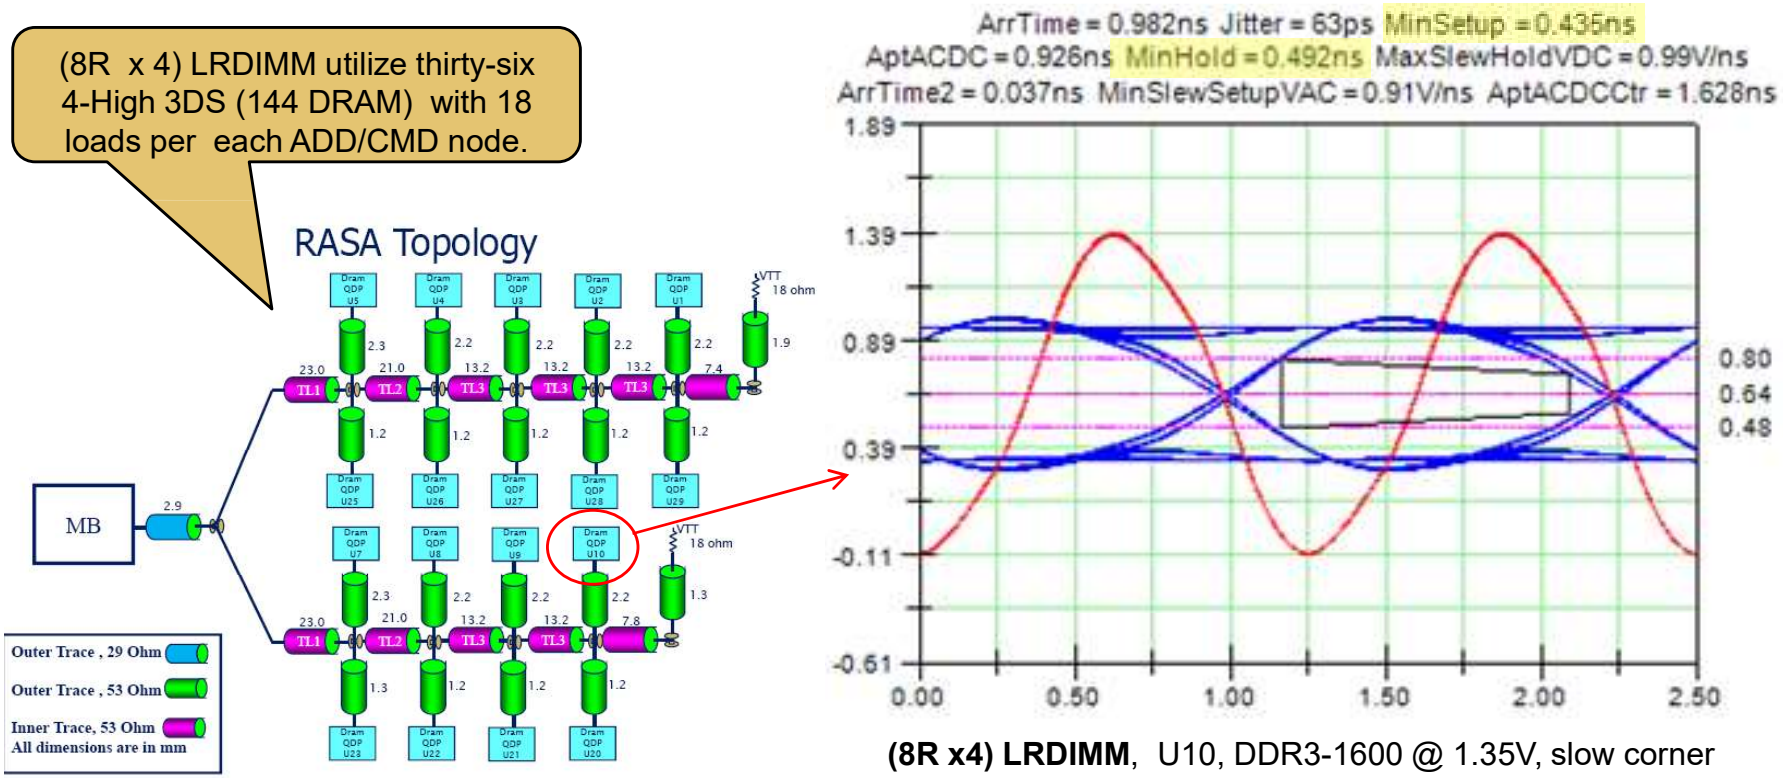

(8R x4) LRDIMM, U10, DDR3-1600 @ 1.35V, slow corner

Potential server applications

- LRDIMM is the most likely server module type that will utilize 3DS devices.
  - LRDIMM support rank multiplication.
  - LRDIMM are will be used to achieve the highest memory possible in the channel, as such they will reflect mostly QR and 8R configurations where stacking is already required thus are good candidates for 3DS memory.

Potential server applications

- A comparison of 3DS RDIMM to LRDIMM utilizing 4Gb based 3DS parts.
  - At this time, RDIMM are limited to four ranks due to registers not supporting rank multiplication.
  - All LRDIMM will have single DQ load per slot, RDIMM will have either a single or dual DQ load per slot depending on architecture.

	(3DS) LRDIMM		(3DS) RDIMM	
	2-High	4-High	2-High	4-High
DR x4	16GB	NA	16GB	NA
QR x4	32GB	NA	32GB	32GB
8R x8	32GB	NA	NA	NA
8R x4	NA	64GB	NA	NA

Thank You